Sedano: A News Stream Processor for Business

Ugo Scaiella, Giacomo Berardi, Giuliano Mega, Roberto Santoro
Spaziodati, Italy
{scaiella,berardi,mega,santoro}@spaziodati.eu

arXiv:1608.06876v1 [cs.IR] 24 Aug 2016

ABSTRACT

We present SEDANO, a system for processing and indexing a continuous stream of business-related news. SEDANO defines pipelines whose stages analyze and enrich news items (e.g., newspaper articles and press releases). News data coming from several content sources are stored, processed and then indexed in order to be consumed by Atoka, our business intelligence product. Atoka users can retrieve news about specific companies, filtering according to various facets. SEDANO features both an entity-linking phase, which finds mentions of companies in news, and a classification phase, which classifies news according to a set of business events. Its flexible architecture allows SEDANO to be deployed on commodity machines while being scalable and fault-tolerant.

Keywords

business intelligence; news retrieval; entity linking; text classification

1. INTRODUCTION

Nowadays a deluge of news articles and press releases often bewilders readers; this explains why news aggregators have surged in popularity. While news aggregators do a good job for readers in general, they may not be enough if someone is interested in specific business-related news and events, e.g.: which are the news where a particular company is mentioned? Which news recently talked about company events, like lay-offs or mergers and acquisitions? Which are the interesting events regarding a specific company? For this reason we created SEDANO, a system that is able to ingest, process and enrich news stream in a scalable manner. The news are indexed and offered via a RESTful API, in order to be consumed by our end-user product Atoka[1], a semantic tool for *lead generation*, but also by other business clients that integrate the information provided by SEDANO in their products.

[1]http://atoka.io

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR '16 July 17-21, 2016, Pisa, Italy

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4069-4/16/07.

DOI: http://dx.doi.org/10.1145/2911451.2926730

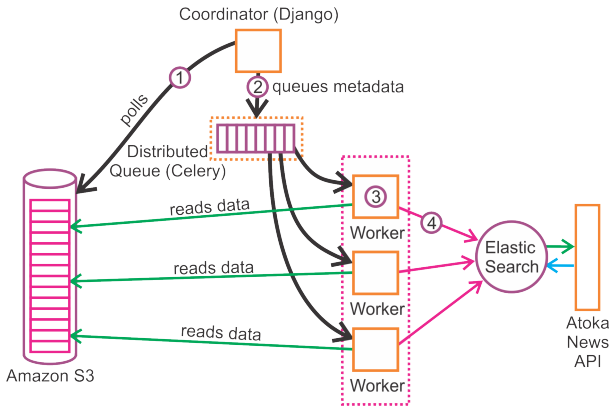

Figure 1: Sedano's base architecture.

2. ARCHITECTURE

SEDANO is based on a simple distributed architecture, depicted in Figure 1. It leverages Amazon's S3[2] for high aggregated throughput, and the Celery[3] distributed queue for control messages. Data is stored in S3 as a collection of *data chunks*, i.e., S3 objects containing *news items*. Processing in SEDANO is triggered by a *coordinator node*, which (1) periodically (i.e., every minute) polls S3 for new chunks and, (2) enqueues *work orders* referring to chunks that require processing into Celery. Once into Celery, work orders become visible to *worker nodes* who then dequeue them, (3) download the corresponding chunks in parallel from S3, and process them independently. At the end of the processing, the resulting news items are (4) indexed into an ElasticSearch[4] cluster. Basic fault tolerance is achieved by having Amazon's Elastic Container Service (ECS)[5] supervise individual components, respawning them into available machines in the event of crashes.

Processing at worker nodes takes place as a series of sequential stages that make up SEDANO's *processing pipelines*. Pipelines are composed of three major steps, discussed next.

Cleansing and normalization. The first step at every pipeline deals with content cleansing and normalization—we extract the relevant content from each news item by processing the HTML in which they are represented, and compute several measures (e.g. the length of an article or if its con-

[2]http://aws.amazon.com/s3
[3]http://www.celeryproject.org
[4]http://www.elastic.co
[5]https://aws.amazon.com/ecs

tent is well-structured) which later help us understand if that news item is "good" or not (and therefore whether or not it should be served to certain clients).

Deduplication. News items are then deduplicated. Since finding duplicates by comparing new documents with previously indexed ones would be costly, we adopt a more scalable on-line approach, based on a novel Locality-Sensitive Hashing (LSH) algorithm which mixes the Solr[6] approach with the Nilsimsa algorithm[7] to hash similar news items into the same code. A news item is then tagged as a duplicate if and only if another news item with the same hash code is known to have been processed before. To keep this information consistent across worker nodes, we rely on a scalable key-value store with strong consistency guarantees.

Enrichment. The remainder of this section discusses the two main enrichment components in the SEDANO pipelines: the *Dandelion Company API* (Section 2.1), an entity linking platform focused on the identification of company mentions in news, and *Selino* (Section 2.2), an ensemble of linear multi-label text classifiers whose classes are business events of interest.

2.1 Dandelion Company API

Dandelion API[8] is a platform for text-analytics as a service: it is the evolution of TAGME [1], a state-of-the-art entity linking system based on a knowledge graph extracted from Wikipedia, whose main benefits are its performance on short texts and its speed. Our partnership with Cerved[9], the leader of business information in Italy, gave us access to data of all Italian companies. We used this data to specialize the entity linking system into the Dandelion Company API for identifying company mentions. It includes information on almost 3M companies, which are added as new entities to our knowledge graph. Despite the huge number of entities that have been added, the system preserves its speed.

Since the new entities do not have cleaned data like Wikipedia ones, they may introduce noise. For this reason we added a new layer, based on a Named Entity Recognition (NER) classifier, which is able to identify named entities of type "company". The result of the NER classifier is used as an additional feature for the subsequent disambiguation step[10], in order to better identify company entities.

We carried out several tests with data sets created by means of crowd-sourced annotations, and we yield a F_1 of 60%. However, for each annotation we also compute a confidence score; setting a threshold for this score can be used to balance prevision vs. recall. Since for this kind of product we are mainly interested in precision, we decided to set the threshold to maximize the $F_{0.5}$, yielding a precision of $\sim 80\%$ and a recall of $\sim 45\%$.

2.2 Selino

Selino is a component of the SEDANO pipeline dedicated to the automatic classification of news according to a predefined set of categories. The categories we are interested in are business events with negative connotation (i.e.: layoffs, strikes, shutdowns, material damages, financial losses, frauds and legal issues) plus other generic events (mergers and acquisitions, new product launches and management changes).

We first built a training set for each category. Given the small size of the taxonomy we created an independent data set for each category; we chose CrowdFlower[11] as crowdsourcing platform for the manual annotation of news. Instead of randomly sampling news for each annotation job, we selected sets of articles ($\sim 12K$ samples per category) through a boolean retrieval model. Each category query was made by expanding a seed query of event related terms. The expansion was performed by (a) creating a vector space model of Italian word semantics using Word2Vec [2], (b) identifying the most similar words to the seeds, (c) manually selecting words for defining the query. This approach allowed us to minimize the manual annotation work and to obtain balanced training sets (each one is made of $\sim 30\%$ positive samples).

Our approach to event detection is through supervised learning: we chose a logistic regression algorithm as text classifier. This linear model has proven to be effective for text classification [3] and it is also reasonably fast. We evaluated the classifiers through 3-fold cross validation, obtaining good results (i.e. an average precision of 0.8 and an average recall of 0.6), especially on precision. The latter is our critical evaluation measure, since a false positive on a negative event can produce unpleasant image damages to the companies cited in an article.

3. CONCLUSIONS

SEDANO is an effective system for processing news and extracting rich information about companies and interesting signals for the business domain. We have shown its design and components, among which the main ones are the entity linking system for identifying company mentions and the events detection classifier. SEDANO has proven to be reliable and scalable, analyzing and indexing hundreds of thousands of Italian news items every day.

We plan to improve further this application by:

- Integrating new components, such as a clustering algorithm to group together related news.
- Moving from the current static event-detection models to an on-line active learning approach, in which we sample news with "ambiguous" classifications and we annotate them through CrowdFlower.
- Extracting specific facts about events from textual contents, employing more sophisticated event detection techniques.

4. REFERENCES

[1] P. Ferragina and U. Scaiella. Tagme: on-the-fly annotation of short text fragments (by wikipedia entities). In *CIKM*, pp 1625–1628, 2010.

[2] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. Distributed Representations of Words and Phrases and their Compositionality. In *NIPS*, pp 3111–3119, 2013.

[3] T. Zhang and F.J. Oles. Text Categorization Based on Regularized Linear Classification Methods. *Information Retrieval*, 4(1):5-31, 2001.

[6] http://lucene.apache.org/solr

[7] https://en.wikipedia.org/wiki/Nilsimsa_Hash

[8] https://dandelion.eu

[9] https://www.cerved.com

[10] Refer to [1] for further details and terminology

[11] http://www.crowdflower.com